(2)

Naval Research Laboratory

Washington, DC 20375-5320

DTIC
S ELECTE JUN 2 3 1993 D
C

AD-A266 087

NRL/MR/8143--93-7337

Asynchronous Transfer Mode (ATM) Technology: An Overview

JAMES R. WEBSTER

Command Control Communications Computers & Intelligence Branch
Space Systems Development Department

June 4, 1993

93 6 23 046

93-14016

Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave Blank*)

2. REPORT DATE

June 4, 1993

3. REPORT TYPE AND DATES COVERED

Final

4. TITLE AND SUBTITLE

Asynchronous Transfer Mode (ATM) Technology: An Overview

5. FUNDING NUMBERS

6. AUTHOR(S)

James R. Webster

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER

NRL/MR/8143--93-7337

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (*Maximum 200 words*)

This paper presents a brief overview of the emerging network protocol standard known as Asynchronous Transfer Mode or ATM. ATM was developed by the telecommunications industry as part of the long-term shift of communications services to the B-ISDN standard. By taking advantage of the highly reliable fiber optic networks available, ATM offers faster and more efficient data communications than existing protocols such as X.25. A rigorous, layered protocol exists that gives ATM the flexibility to carry many different types of traffic, including voice, imagery, and video, as well as data.

14. SUBJECT TERMS

Networks
Protocols
Telecommunications

15. NUMBER OF PAGES

23

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std 239-18
298-102

CONTENTS

1. INTRODUCTION 1

2. BACKGROUND 1
   - 2.1 ATM vs. Existing Technologies 3
   - 2.2 OSI Model 5

3. ASYNCHRONOUS TRANSFER MODE (ATM) 7
   - 3.1 ATM Communications 7
   - 3.2 ATM Networking 9

4. SUMMARY 12

5. REFERENCES 14

6. GLOSSARY—ACRONYMS & DEFINITIONS 16

ASYNCHRONOUS TRANSFER MODE (ATM) TECHNOLOGY: AN OVERVIEW

1. Introduction

This paper is a brief overview of the emerging network protocol standard known as asynchronous transfer mode or ATM. The paper's intended audience is any engineer or manager with a basic understanding of telecommunication networks and protocols. The paper will first present a general history of the evolution of ATM. Next, a detailed discussion will examine the communications and networking mechanisms defined for ATM. The paper will then conclude with a summary of the main points of this paper and a list of references and a glossary of networking terms.

2. Background

The development of asynchronous transfer mode (ATM) technology is a direct result of recent work in the public telecommunications industry to improve and expand the types of services available. Until recently, most public telecommunications networks have been designed for specific purposes such as telephony and cable TV. Today, there is a growing desire to use public networks for applications such as distributed data processing and the transmission of graphic and video images. These require a combination of variable bandwidth demands and real-time delivery characteristics. In the past, the telecommunications industry responded differently to such demands. For example, when the public telephone network was not able to effectively support data services, other specialized networks were developed, e.g. TELNET and MILNET. Finally, many private networks (e.g. cable TV, BITNET, LANs) exist which deploy non-standardized equipment, interfaces, and protocols; thus, they are unable to offer access to other networks and users. By using a common mechanism to transport non-voice services such as data and video over a single network rather than several separate networks, ATM represents a break with these past network philosophies.

The first example of this new approach to telecommunications within the U.S. public telephone network is the integrated services digital network (ISDN) [1]. ISDN transports voice and digital data communications over a single network. All source information (irrespective of its meaning to the user) is transmitted and switched as digital signals end to end. This allows a common signal transfer mechanism in the network to serve very different applications, e.g., voice and data transmission.

As a further enhancement to ISDN, broadband ISDN (B-ISDN) has been defined to provide considerably higher bit rates and multirate

Manuscript approved March 31, 1993

capability [2]. A key capability of B-ISDN is to provide users the flexibility to choose bit rates on a per connection basis. In the past several years, there has been a significant consensus in both national and international standards bodies in laying the foundation for B-ISDN.

The basic transmission mechanism for B-ISDN is ATM. ATM performs universal transport processing, i.e., multiplexing and routing [3], for B-ISDN and offers several major benefits to network users:

- **Bandwidth Efficiency:** ATM is a packet-switching technology that makes efficient use of communications channels. All users within a multiplexed channel are granted bandwidth only when they require network services. Idle users do not waste channel resources with empty packets or frames.

- **Scalability:** The ATM transport mechanism can scale its speed to suit different network configurations, allowing data to be transferred from one network to another in a common format.

- **Transparency**: ATM allows free mixture of different types of data, such as voice and digital data, within the same connection. The network circuits established are transparent to the type of data being transported.

- **Granularity:** The existing public telephone network is limited to only a few choices (DS0, DS1, or DS3) for channel rates. ATM allows the user to choose the specific rate desired for the application.

- **Network Flexibility:** In the future, ATM will simplify interconnecting LANs and wide-area networks (WANs). The interface will require little more than address manipulation, in contrast to the complex protocol conversion techniques currently necessary.

ATM combines the advantages of both circuit-switched and packet-switched techniques. Like circuit-switched networks, ATM transfers data sequentially, which simplifies the user to network interface (UNI) because buffering of packets is not required. It is also vital for communication applications such as voice where sequential delivery of data is essential. ATM is a packet-switching technology, however. All data is packetized at

the UNI and transferred through the network from node to node (between NNIs or network-node interfaces) along a virtual circuit until it reaches the destination. The use of the UNI and NNI in an ATM network is depicted in figure 1. A virtual circuit is a *logical* connection between two users, where the delivery of data does not depend on a single, unique path through the network. Because packet-switched networks use virtual circuits as opposed to the permanent circuits of circuit-switched networks, the network can establish connections with much more flexibility. The actual methods for establishing and using an ATM connection will be discussed in detail in section 3 of this paper.

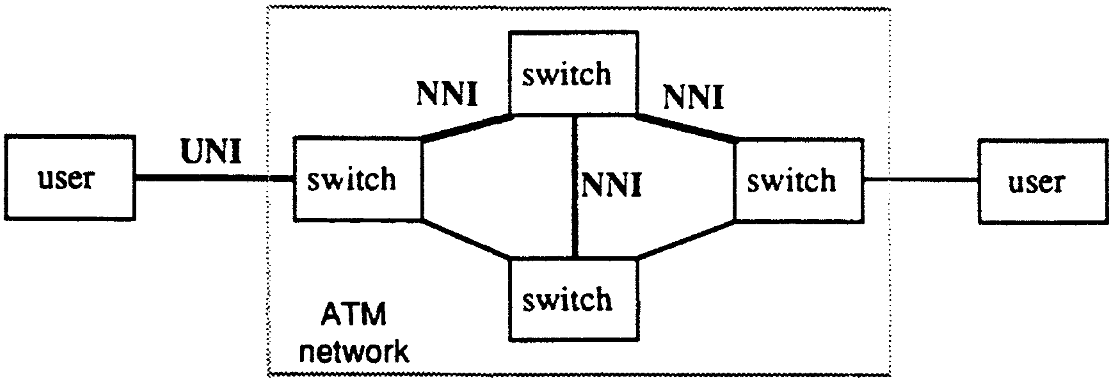

FIGURE 1. General ATM Network Model

The current recommendations for B-ISDN include the Synchronous Optical NETwork (SONET), also known as Synchronous Digital Hierarchy (SDH), as the physical-layer transmission standard [4]. ATM, together with SONET, is expected to provide the reliable high-speed transport, bandwidth flexibility, and integrated transmission and switching for a diverse set of traffic characteristics as required by B-ISDN.

Many network technologies exist in industry and it is important to contrast their capabilities with those of ATM. The following discussion examines the main advantages of using an ATM network approach over current WAN and local-area network (LAN) technologies.

2.1 ATM vs. Existing Technologies

Most existing WAN protocols were developed for use over noisy, error-prone analog transmission facilities which required very robust error correction capabilities. As a result, these protocols were inherently

slow. For example, X.25 uses a rigorous procedure to check the integrity of the message structure at every node in the channel [5]. Finally, before data from LAN can be carried over a WAN, or vice-versa, it must first undergo a time-consuming protocol conversion process. Even with today's high-speed processing and decreasing memory costs, protocol conversion speeds will continue to lag behind emerging requirements for network services. By using a much simpler error-control mechanism and a common data format across LANs and WANs, ATM removes many of the shortcomings of X.25-style WANs.

The primary drawback with using local area networks such as Token Ring, Fiber-Distributed Data Interface (FDDI) and Ethernet is that all network users have equal access to a fixed LAN bandwidth. As more users are added, contention for the available bandwidth increases proportionally. Upgrades to Ethernet and FDDI have been proposed to address this problem: the Ethernet switch [6] and FDDI-II [7] respectively. Recently, support for these solutions has diminished due to the increasing interest in ATM. Many telecommunication vendors are redirecting resources towards developing ATM technologies because, along with its support for real-time applications, ATM can supply high-capacity bandwidth more efficiently [8].

When bandwidth is set aside for specific services, as in FDDI-II, the amount of bandwidth allocated must respond to the peak burst rate. If 12 Mbps is required for a peak burst, then a 12 Mbps channel must be allocated for the duration of the session. If more network users start to perform tasks requiring high data rates, bandwidth requirements escalate and frame delivery becomes less predictable. ATM's ability to deliver bandwidth dynamically in response to users' needs prevents this kind of problem.

Frame relay, a fairly recent development, is a protocol that was specifically designed to transfer LAN traffic over telecommunications circuits. It takes advantage of the growing number of high quality, fiber-based communication networks and their virtually error-free performance. As a result, most of X.25's burdensome checking can be eliminated in favor of one or two simple checks: address validity and frame integrity. Frames which fail either check are discarded, leaving the processes at the ends of the channel to recover from the loss of the frame.

Frame relay accepts LAN traffic, which commonly generates variable-length frames of data, and adds only a wide area network address

at the front and its own check sequence at the end. However, because frame relay is so well suited to LAN-style data communications, it is not flexible enough to cope with future heterogeneous network traffic: data, voice, image, and video. While data is well served by frame relay's variable-length format, services like voice will suffer if not transmitted in a reasonably consistent manner.

For example, a T1 channel transmits a voice sample once every 125 microseconds with good accuracy. If the interval between samples or groups of samples varies too much, a problem called "jitter" arises when the signal is reconstructed. If intervals become very large, echoes result. If voice samples are sandwiched between variable-length data frames, as in frame relay, jitter and echo can arise easily. ATM uses short frames, called "cells," of fixed length which minimize these problems. Because samples now arrive at regularly spaced intervals, network latency, or end-to-end delay which can cause jitter and echoes, is minimized through the network.

Understanding the reasons why cell-switched, broadband networks, and specifically ATM, are able to perform at higher speeds and with more flexibility than existing networks requires an appreciation of the ATM protocol in context of the International Standards Organization (ISO) Open Systems Interconnect (OSI) model [9].

2.2 OSI Model

The OSI model is of fundamental importance in the study of data and computer communications. It provides an architecture within which protocol standards can be developed and a common frame of reference for the purpose of discussion. As defined by the CCITT (the international telecommunications standards body), this model consists of seven layers:

7 - Application	Provides specific services for applications such as file transfer
6 - Presentation	Provides data translation between systems
5 - Session	Provides synchronization of data flow
4 - Transport	Provides end-to-end data transmission integrity and error recovery

3 - Network Switches and routes information units between local and remote nodes

2 - Link Corrects transmission errors and provides flow control between adjacent nodes

1 - Physical Transmits bit stream on physical medium

ATM is based on the layered architecture concept of the OSI model. Each layer provides services to the layer above and the aggregate effect enables communications between top layer processes, such as applications at the ends of a connection. An important concept behind layering is the ability to revise or change a layer without impacting the layers above or below. Thus, the physical layer protocol for ATM may be changed with no impact on the ATM layer above, or the services the ATM layer provides to higher layers.

ATM is contained within the B-ISDN protocol stack as shown in figure 2. The ATM-relevant portions comprise the bottom three layers: the ATM adaptation layer (AAL) [10], divided into two sublayers; the ATM layer [11]; and the physical layer [12], also divided into two sublayers. The higher layers within the B-ISDN protocol stack have not been undefined. There is not a one-to-one correspondence between the B-ISDN and OSI stack however. Even though it contains three layers, the ATM portion of the B-ISDN protocol stack is roughly equivalent to OSI Layer 1 and part of Layer 2.

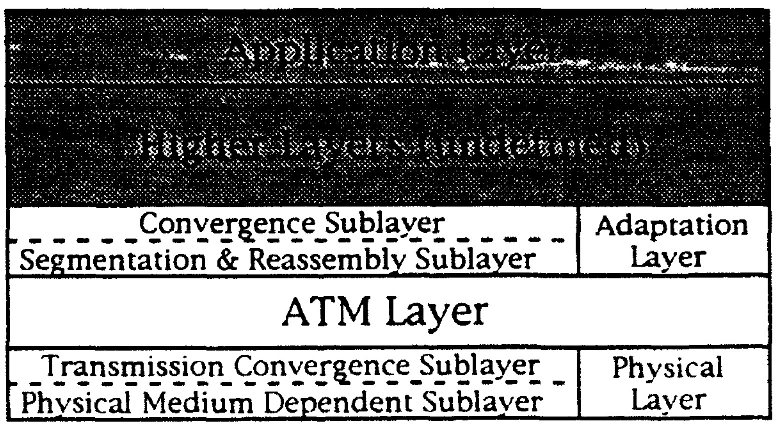

FIGURE 2. B-ISDN Protocol Stack

The B-ISDN physical layer defines the medium for transmission and any medium dependent parameters such as rates and levels, as well as the method, such as framing, for finding the data contained within the

medium. While copper media, such as coaxial cable, could be used at the lower rates proposed for B-ISDN (e.g., 155 Mbps), higher rates would demand optical fiber because of its inherently lower error rates [13].

Like other OSI Layer 2 processes, the ATM layer is concerned with the data transmission between two adjacent network nodes (link-by-link), usually not the end points of a connection (end-to-end). Thus, cell addressing is of local significance only - between adjacent nodes or switches (i.e., NNIs). The ATM layer also provides for the basic 53-byte cell format and defines the header contents.

The AAL adapts higher level data into formats compatible with the ATM layer requirements. This layer is dependent on the higher layer services being transported. Currently, several different AALs have been defined for services such as data-only transport, voice, video, and others. The AAL is an end-to-end process used only by communicating entities to insert and remove data from the ATM layer. Once the virtual circuits are established, all processing of cells occurs at OSI Layer 2 as the cells follow the same path. This is much more efficient than OSI Layer 3 routing, and has the advantage that ATM switching [14] can be implemented in silicon, minimizing latency and increasing throughput.

With this brief introduction to the evolution of ATM we can now turn our attention to the basic mechanisms of the ATM protocol.

3. Asynchronous Transfer Mode (ATM)

3.1 ATM Communications

The bottom three layers of the B-ISDN stack define the ATM communications process. However, the bulk of the message processing is performed at the middle, or ATM, layer. At this layer, the fundamental unit of communications is formed: the ATM cell. It contains the information to be transferred and all the information the network needs to relay the cell from one node to the next over an ATM connection.

ATM cells consist of two parts: a 5-octet header and a 48-octet information payload as shown in figure 3 (the octet is the international equivalent of a byte). For networking purposes, only the header is significant. The AAL layer divides higher layer data units (created by user applications) into 48-octet segments. These data segments are then passed

to the ATM layer and prefixed with an initial header for routing and error control purposes. This header consists of the following fields:

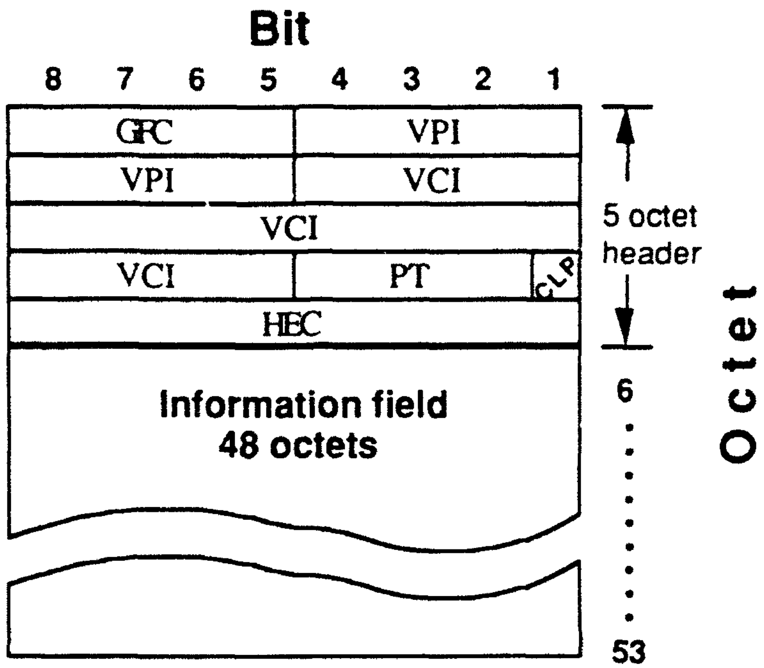

FIGURE 3. ATM Cell Structure

GFC (Generic Flow Control)
The first four bits of the first octet comprise a generic flow control (GFC) field. It is used to control the flow of traffic across the user-network interface (UNI), e.g., from workstations to a network switch. This field is used only at the UNI; between network-node interfaces (NNI), such as between switches, these four bits provide additional network address capacity. Exact mechanisms for flow control are still being developed, and no explicit definition for use of this field exists at this time.

VPI/VCI (Virtual Path/Channel Identifier)
The next 24 bits constitute the ATM address. This three-octet field is divided into two subfields. The first octet contains the virtual path identifier (VPI) and the second two octets make up the virtual channel identifier (VCI). Virtual paths and channels are discussed in a later section on ATM networking; for now the entire field may be thought of as providing an OSI Layer 2 address.

PT (Payload Type)
The next three bits, PT or payload type, indicate the type of information carried by the cell. ATM cells will be used to carry different types of user information that may require different handling by the network or terminating equipment. Cells will also be used to transfer operations and maintenance messages across the network between users or between a user and a service provider. Three bit binary values will indicate the type of message in the payload. At this time, values 0-3 are reserved for identifying various types of user data, 4 and 5 indicate management information, 6 and 7 are reserved for future definition.

CLP (Cell Loss Priority)
The last bit of octet four, CLP, indicates the cell loss priority, and is set by the user. This bit indicates the eligibility of the cell to be discarded by the network under congested conditions.

HEC (Header Error Control)
The final octet, the HEC, is the header error control field. This is a cyclic redundancy check (CRC) error-correcting code calculated across the previous four octets of the header. It primarily provides protection against misdelivery of cells due to address errors and also detects multiple header errors and corrects single bit errors. Note: the HEC does not provide any indication of the quality of data in the information field.

Information Field
Following the HEC is the 48-octet information field containing the user data. Inserting user data into the information field is accomplished by the ATM adaptation layer (AAL). It should be noted that depending on the type of adaptation process, not all 48 octets are user information [15]. Up to four octets may be used by the adaptation process itself.

3.2 ATM Networking

ATM is asynchronous in the sense that in the context of the underlying physical media, cells allocated to the same connection may exhibit an irregular recurrence pattern. This is in contrast to *Synchronous* Transfer Mode (STM), where a data unit associated with a given channel is

identified by its position in the transmission frame. These concepts are illustrated in figure 4.

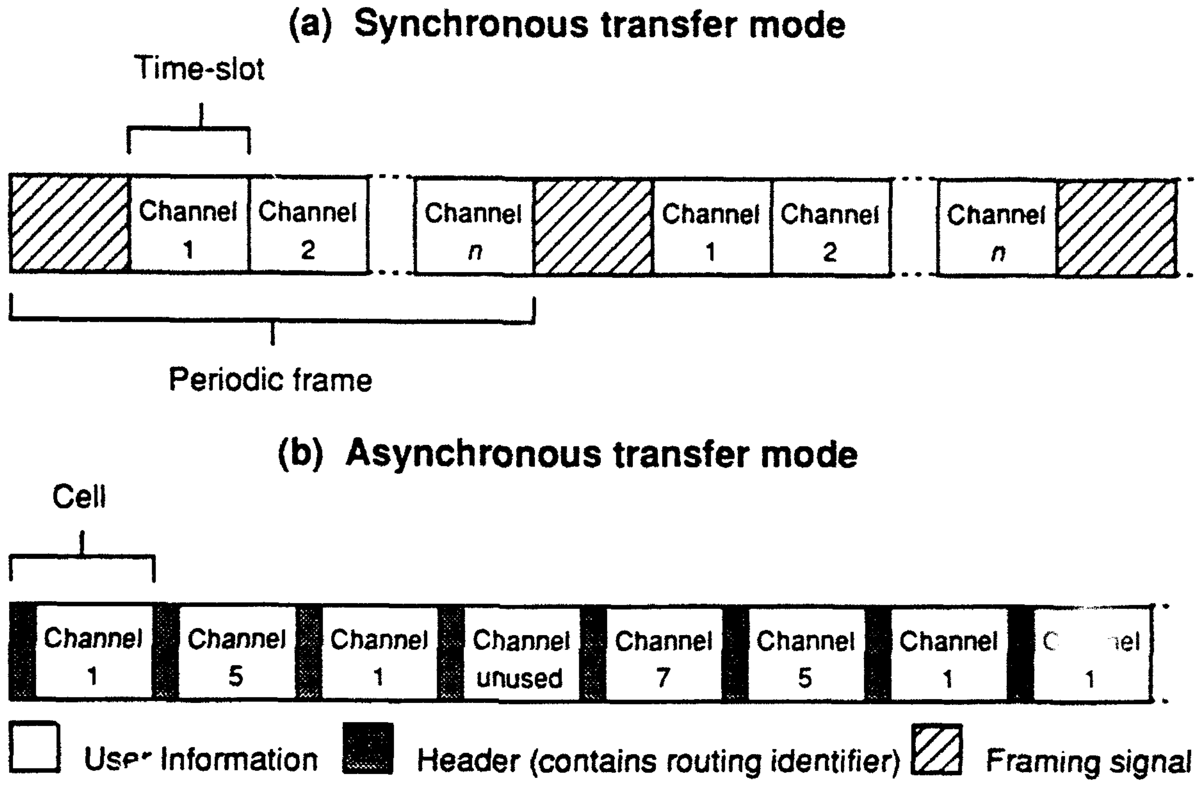

FIGURE 4. STM and ATM Principles

In ATM-based networks the multiplexing and switching of cells is independent of the actual application. Therefore, in principle, the same equipment can handle a low bit rate connection concurrently with a high bit rate connection comprised of data in streams or bursts. STM implementations such as the synchronous time division multiplexing (TDM) approach used for narrowband ISDN [16] provide only fixed-rate channels. Many applications of the types suggested by B-ISDN (such as multiplexed voice and video) cannot be accommodated by this approach, due to the inflexibility of the TDM structure to rate adaptation and the difficulty in switching data streams of multiple high data rates. These data applications are bursty in nature and are well suited to the packet-switching approach of ATM.

TDM identifies the individual users utilizing the communications channel by their position in time, or "slot," relative to the framing bit. A

user can always be found in a specific slot synchronized to the frame bit, hence, Synchronous Transfer Mode. Thus, bandwidth is allocated even when no data is being transmitted, which reduces efficiency. In contrast, ATM utilizes *labeled*-channel multiplexing. This allows ATM users to remain idle, but still connected, without wasting bandwidth. Label multiplexing uses a label in the cell header, called the "connection identifier", to explicitly define the virtual channel on a physical link to which the cell belongs. The connection identifier consists of two sub-fields, the virtual channel identifier (VCI) and the virtual path identifier (VPI) mentioned previously.

ATM implements the concept of virtual channels (VCs) and virtual paths (VPs) to accomplish network routing. A virtual channel is a connection between two communicating ATM entities. It may consist, for example, of a concatenation of several ATM links - a workstation to a central switch, a switch to a SONET link. All communications proceed along this same VC which preserves cell sequence and provides a certain grade or quality of service.

A virtual path is a group of VCs carried between two points and may also involve many ATM links. While the VCs are associated with a VP, they are not unbundled or processed in any way. The cell sequence of each VC is still preserved and the grade of service of the VP is established by the most demanding of the constituent VCs.

It is important to remember that the cell header contains both a VPI and a VCI, allowing a cell to be given a unique VC identifier and be associated with a particular virtual path by having its VPI in common with other VCs. A cell may also not be associated with any VP, in which case it would have a null VPI and only a unique VCI. Thus, the same VCIs may be used unambiguously within different VPs.

VPs provide a convenient method for "bundling" traffic all heading for the same destination. Switching hardware only needs to check the VPI portion of the header to relay the traffic rather than the entire three octet address. VPs can be established by the network for trunking purposes or by users to define the structure of their private networks.

VPIs and VCIs are used to establish virtual connections, either permanent or on demand. These connections are termed "virtual" because they exist in the physical sense only for as long as the message traverses

the ATM element (switch, cross connect, etc.). TDM technology establishes a real, physical, dedicated channel for the duration of a call or as long as a customer desires a private line between two points. That channel consists of physical wires, reserved time slots in a TDM multiplexer, physical cross point connections in a switch, or other real connections. For the duration of the call, no other traffic may share the dedicated facility, simply because there is no way to identify separate traffic in the channel. The only identification is the time slot number or other physical association between user and channel.

By using OSI Layer 2 addressing, ATM separates the traffic from the physical channel and allows many users to share facilities. Consecutive cells of totally independent users may flow down the same facility and into the same cross-connect port, for example. The cross-connect will identify each cell and route it according to the address, not according to the input port as is the case with TDM. Thus, the communication channel exists only virtually, in the form of look-up tables and associations of input and output addresses.

The idea of virtual circuits is to keep the network fully utilized, thereby conserving resources. Rather than dedicating full time channels to users, network capacity is shared. Today, the communication capacity of a network is often wasted by idle users or limited by high-bandwidth applications. For example, voice communication is taken up by pauses (listening or other quiet intervals), while data communications involve uni-directional transfers of large volumes of data, with only brief acknowledgments flowing the other way. LAN interconnections are even worse with short, bursty transfers followed by silence.

4. Summary

ATM is a network protocol developed by the telecommunications industry as part of the long-term shift in telecommunications services to B-ISDN. By taking advantage of the highly reliable fiber optic networks available, ATM offers faster and more efficient data communications than previous protocols such as X.25. A rigorous, layered protocol exists that gives ATM the flexibility to carry many different kinds of traffic, including voice, image, and video, as well as data. The layered protocol also allows ATM cells to be transported over a wide variety of media in both public

and private networks, providing superior compatibility between local and wide area networks.

5. References

[1] Stallings, W.: *ISDN and Broadband ISDN*. MacMillan Publishing Company, New York, 1992.

[2] Handel, R. and Huber, M. N.: *Integrated Broadband Networks - An Introduction to ATM-Based Networks*. Addison-Wesley, New York, 1991.

[3] Kawarasaki, M.: B-ISDN Architecture and Protocol. *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 9, December 1991, pp. 1405 - 1415.

[4] Cheung, N. K.: The Infrastructure for Gigabit Computer Networks. *IEEE Communications Magazine*, April 1992, pp. 60 - 68.

[5] Stallings, W.: *Handbook of Computer-Communications Standards, Volume 1: The Open Systems Interconnections (OSI) Model and OSI-Related Standards, Second Edition*. Macmillan, New York, 1990.

[6] Lamb, C.: Speeding to the ATM. *Unix Review*, vol.10, no. 10, pp. 29 - 36.

[7] ANSI: ANSI Standard 'Hybrid Ring Control'. Revision 6, May 1990.

[8] Raymond, M.: Real-Time Delivery Could Revolutionize Networking. *digital news & review*, September 28, 1992, pp. 8 - 12.

[9] CCITT: Recommendation X. 200. 'Reference Model of Open System Interconnection for CCITT Applications'. *Blue Book*, Fascile VIII.4, Geneva, 1989.

[10] CCITT: Recommendation I.363. 'B-ISDN ATM Adaptation Layer (AAL) Specification', Geneva, 1991.

[11] CCITT: Recommendation I.361. 'B-ISDN ATM Layer Specification', Geneva, 1991.

[12] CCITT: Recommendation I.432. 'B-ISDN User-Network Interface - Physical Layer Specification', Geneva, 1991.

[13] Ritchie, G. R.: Impact of Fiber Optics on Local Exchange Networks. *Opt. Fiber Commun. Conf. Digest*, paper ThH3, San Jose, CA, Feb. 2-7, 1992, p.225.

[14] Newman, P.: ATM Technology for Corporate Networks. *IEEE Communications Magazine*. April 1992, pp. 90 - 101.

[15] Lane, J.: *Asynchronous Transfer Mode: Bandwidth for the Future*. Telco Systems, 1992.

[16] Bocker, P.: *ISDN - The Integrated Services Digital Network - Concepts, Methods, Systems*. Spinger-Verlag, New York, 1988.

6. Glossary - Acronyms & Definitions

AAL
ATM Adaptation Layer

ANSI
American National Standards Institute - its T1 committee is charged with developing US telecommunications standards.

Asynchronous
Two or more signals sourced from independent clocks, therefore having different frequency and phase relations.

ATM
Asynchronous Transfer Mode - a transport and switching method in which information does not occur periodically with respect to some reference such as a frame pattern.

B-ISDN
Broadband Integrated Services Digital Network - a common digital network suitable for voice, video, and high-speed data services running at rates beginning at 155 Mbps.

CCITT
International Telegraph and Telephone Consultative Committee - the international standards body for telecommunications.

Circuit-Switching
A method of communicating in which a dedicated communications path is established between two devices through one or more switching nodes. Unlike packet-switching, digital data is sent as a continuous stream of bits. Data rate is guaranteed, and delay is essentially limited to propagation time.

Connectionless Service
A type of service in which no predetermined path or link has been established for transfer of information.

Connection-Oriented Service

A type of service in which information always traverses the same pre-established path or link between two points.

CRC
Cyclic Redundancy Check - an error detection and correction algorithm.

DS0, DS1, DS3
Electrical signal specifications for specific data rates within the DS-n asynchronous digital hierarchy of TDM carriers:

DS0	64 kbps	1	Voice Channel
DS1	1.544 Mbps	24	DS0s
DS1C	3.152 Mbps	2	DS1s
DS2	6.312 Mbps	4	DS1s
DS3	44.736 Mbps	28	DS1s
DS4	274.176 Mbps	168	DS1s

Frame
A specific pattern of bits used to identify the beginning or end of a frame.

Frame Relay
A recently developed, fast packet switching protocol based on the LAPD protocol of ISDN that performs routing and transfer with less processing than X.25.

HEC
Header Error Control - a CRC code located in the last octet of an ATM cell header used for checking header integrity only.

ISDN
Integrated Services Digital Network - a worldwide telecommunication service that will use digital transmission and switching technology to support voice and digital data communications.

Isochronous
Signals carrying embedded timing information or dependent on uniform timing. Data has no embedded timing - the information contained is not time dependent. Voice and video are intimately tied to time. With TDM services there is a direct relationship between the signal rate used to digitize the voice samples and the bearer channel rate, allowing accurate

reconstruction of signals, such as voice, at the receiving end. Services like ATM must use care in transferring such signals so timing can be recovered, since it cannot be derived from the ATM channel itself.

Nibble
4 bits (half a byte or octet).

NNI
Network-Node Interface - the interface between two public network pieces of equipment, such as switches.

Octet
8 bits (synonymous to a byte).

OSI
Open Systems Interconnect - a model for data communications which allocates well defined functional requirements into seven distinct layers.

Packet-Switching
A method of transmitting messages through a communications network, in which long messages are subdivided into short packets of data. Each packet is passed from source to destination through intermediate nodes. At each node, the entire message is received, stored briefly, and then passed on to the next node.

PDU
Protocol Data Unit - a segment of data generated by a specific layer of a protocol stack; usually contains information from the next higher layer encapsulated with header and trailer data generated by the layer in question.

PVC
Permanent Virtual Channel (or Circuit) - a channel through an ATM network provisioned by a carrier between two end points, used for dedicated long term information exchange between locations.

SDH
Synchronous Digital Hierarchy - see SONET.

SONET

Synchronous Optical NETwork - a new and growing body of standards that defines all aspects of transporting and managing digital traffic over fiber optic facilities in the public network.

STM
Synchronous Transfer Mode - A transport and switching method that depends on information occurring in regular and fixed patterns with respect to a reference such as a frame pattern.

SVC
Switched Virtual Channel (or Circuit) - a channel established on demand by network signaling used for information transport between two locations, lasting only for the duration of the transfer; the datacom equivalent of a dialed phone call.

Synchronous
Signals that are sourced from the same timing reference and hence are identical in frequency.

TDM
Time Division Multiplexing - traditional digital multiplexing where a signal occupies a fixed, repetitive time slot within a higher-rate signal.

UNI
User-Network Interface - the physical and electrical demarcation point between the user and the public network service provider (switch, etc.).

VBR
Variable Bit Rate - referring to processes such as LANs which generate messages in a random, bursty manner rather than continuously.

VC
Virtual Channel - a communications path between two nodes identified by label rather than fixed physical path.

VCI
Virtual Channel Identifier - the address or label of a VC.

Virtual Circuit

A packet-switching mechanism in which a logical connection (virtual circuit) is established between two users at the start of a transmission. All packets follow the same route, need not carry the same address, and arrive in sequence.

VP

Virtual Path - a collection of VCs all traveling between common points.

VPI

Virtual Path Identifier - the address of a Virtual Path.

X.25

a well-established data switching and transport method that relies on a significant amount of processing to ensure reliable transport over metallic facilities and media.